

CLAIMS

I claim:

1. A method for automatic generation of join graphs for relational database queries, comprising:

(a) receiving an input list of tables including attributes of interest for a database query;

(b) marking instances of tables of said input list having single occurrences in an hierarchical representation of a database schema, and marking ancestors of said instances of tables according to said hierarchical representation;

(c) marking unmarked instances of multi-dimensional tables of said input list closest to marked instances, marking unmarked ancestors of said unmarked instances of said multi-dimensional tables according to said hierarchical representation, and marking unmarked instances of one-dimensional tables that reference said multi-dimensional tables and have said unmarked instances of said multi-dimensional tables as parents according to said hierarchical representation;

(d) marking unmarked instances of one-dimensional tables of said list of tables closest to marked instances, and marking unmarked ancestors of said unmarked instances of said one-dimensional tables according to said hierarchical representation; and

(e) generating a join graph corresponding to said input list of tables from said marked instances in said hierarchical representation.

2. The method according to claim 1, wherein said (b) comprises:

(b1) selecting a table of said input list;

(b2) determining whether a number of instances of said table in an hierarchical representation of a database schema is greater than one;

(b3) if said number of instances of said table is not greater than one, then marking an instance of said table in said hierarchical representation and all unmarked instances of ancestor tables to said table according to said hierarchical representation; and

(b4) repeating (b1) to (b3) until a last table in said input list is processed.

3. The method according to claim 1, wherein said (b) comprises:

(b5) selecting a table having a table name from said input list;

(b6) determining whether a number of dimensions of said table is greater than one according to a database schema;

(b7) if said number of dimensions of said table is greater than one, then adding said table name to a multi-dimensional list of tables, and if said number of dimensions of said table is not greater than one, then adding said table name to a one-dimensional list of tables; and

(b8) repeating (b5) to (b7) until a last table in said input list is processed.

4. The method according to claim 1, wherein said (b) comprises:

(b1) selecting a table having a table name from said input list;

(b2) determining whether a number of instances of said table in an hierarchical representation of a database schema is one;

(b3) if said number of instances of said table is one, then marking an instance of said table in said hierarchical representation and all unmarked instances of ancestor tables to said table according to said hierarchical representation;

(b4) if said number of instances of said table is greater than one, then determining whether a number of dimensions of said table is greater than one according to a database schema;

(b5) if said number of dimensions of said table is greater than one, then adding said table name to a multi-dimensional list of tables, and if said number of dimensions of said table is not greater than one, then adding said table name to a one-dimensional list of tables; and

(b6) repeating (b1) to (b5) for all tables in said input list.

5. The method according to claim 3, wherein said (c) comprises:

(c1) selecting a multi-dimensional table from said multi-dimensional list;

(c2) if an instance of said multi-dimensional table has been marked in said hierarchical representation,

then selecting another multi-dimensional table from said multi-dimensional list;

(c3) if an instance of said multi-dimensional table has not been marked in said hierarchical representation, then finding an instance of said multi-dimensional table that is a closest child in relationship to a marked instance in said hierarchical representation;

(c4) marking said found instance of said multi-dimensional table and instances of all unmarked ancestors of said found instance of said multi-dimensional table according to said hierarchical representation;

(c5) generating a dimensions list for said multi-dimensional table;

(c6) processing said dimensions list; and

(c7) repeating (c1) to (c6) until a last table in said multi-dimensional list is processed.

6. The method according to claim 5, wherein said (c5) comprises referring to a database schema to determine all reference tables to said multi-dimensional table, and including all said reference tables except a parent of said multi-dimensional table according to said hierarchical representation.

7. The method according to claim 6, wherein said (c6) comprises:

(c61) selecting a reference table from said dimensions list;

(c62) if a table name of said reference table is not in said one-dimensional list, then selecting another reference table from said dimensions list;

(c63) if said table name of said reference table is in said one-dimensional list, then marking an instance of said reference table in said hierarchical representation having said current multi-dimensional table as its parent;

(c64) removing said table name from said one-dimensional list; and

(c65) repeating (c61) to (c64) until a last reference table in said dimensions list has been processed.

8. The method according to claim 1, wherein said (d) comprises:

(d1) selecting a one-dimensional table from said one-dimensional list;

(d2) if an instance of said one-dimensional table has been marked in said hierarchical representation, then selecting another one-dimensional table from said one-dimensional list;

(d3) if an instance of said one-dimensional table has not been marked in said hierarchical representation, then finding an instance of said one-dimensional table that is a closest child in relationship to a marked instance in said hierarchical representation;

(d4) marking said found instance of said one-dimensional table and instances of all unmarked ancestors of said found instance of said one-dimensional table according to said hierarchical representation;

(d5) repeating (d1) to (d4) until a last one-dimensional table in said one-dimensional list is processed.

100-434260